BASEMENT:

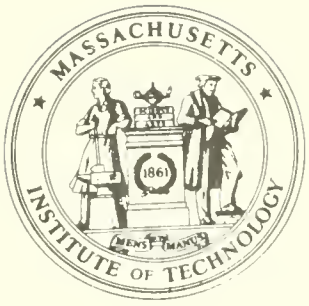

LIBRARY
OF THE
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY

HD28
.M414
no.373-
69

WORKING PAPER
ALFRED P. SLOAN SCHOOL OF MANAGEMENT

TOWARDS ASSOCIATIVE INFORMATION SYSTEMS FOR PLANNING AND CONTROL

by

Kenan E. Sahin*

373-69

Spring, 1968

MASSACHUSETTS
INSTITUTE OF TECHNOLOGY
50 MEMORIAL DRIVE
CAMBRIDGE, MASSACHUSETTS 02139

MASS. INST.
MAY 2 19
DEWEY LIBRARY

TOWARDS ASSOCIATIVE INFORMATION SYSTEMS FOR PLANNING AND CONTROL

by

Kenan E. Sahin*

373-69 Spring, 1968

* Assistant Professor of Management, MIT.

TOWARDS ASSOCIATIVE INFORMATION SYSTEMS FOR PLANNING AND CONTROL

by

Kenan E. Sahin

Introduction

The ideas to be presented and discussed here stem from and form the initial phase of a research into design and implementation of associative management information systems for utilization in planning and control. In this paper we will discuss the basic purpose and the motivation for it; then we will state some assertions with respect to the organization of the data base, review state of the art, present a scheme for an associative information system and finally analyze some useful, intelligent activities that can be achieved using such a system.

Purpose and Motivation

The basic purpose is to impart more "intelligence" to the information system. The motivation for this derives from ever faster increase in the complexity within which a firm has to operate.

We would like to differentiate between two types of complexity. In one case, the relationships are involved and complicated but extend over a small domain (intensive complexity). In the other, the relationships (underlying patterns) are fairly simple but they involve or cover or describe a large body of data (extensive complexity). It is to this latter type of complexity that we shall address ourselves.

The judgmental faculty of human mind is quite remarkable. At one sight, a person can perceive a multitude of relationships, form associations with relevant items in his memory, perform many types of transformations

679311

on data taken in and then generate many hypotheses, decisions, generalizations, and "bright ideas." However, the human being as a decision making unit is handicapped in three respects: (1) his input channels have definite capacities, i.e., conform to Shannon's formulations, (2) although his memory is fairly large, with respect to given problems the data base that he can maintain in his head is limited in size, and (3) while he can with great ease process and transform and link incoming data and data already in memory, the precision of the operations as well as that of his output is not impressive. Contributing greatly to this is the continuous decay in his memory.

These handicaps are most apparent when the decision maker is faced with the extensive complexity (relationships are simple but are involved in and contained by a massive data base) for which only a low level of intelligence but large input capacities, ability to maintain a large data base, and ability to manipulate it with ease are needed. On the other hand, computerized information systems designed with new concepts can exhibit lower levels of intelligence in large scale activities and thus both can allow the firm to deal more effectively with extensive complexity and also leave decision makers more time for tackling problems requiring much more sophistication and ingenuity. Such an organization of the information system could be termed "decentralization of intelligence" in the man-machine system. One way of representing the impact of such an approach is shown below:

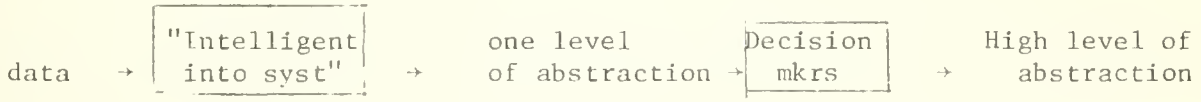

Such reorganizations of the information system can and should be done.

Against this motivational background, some specifics will be presented.

Assertions

We shall post five assertions below and give some justifications. The rest of the paper should provide more justification.

(1) Reorganization of data base. The data base should be organized so as to make "content-addressability" possible, i.e., specifying either all or part of the contents of an image* should be all that is needed to retrieve appropriate images from memory. This contrasts with today's computerized systems relying on "location numbers" or "addresses" for retrieval.

Present day approach derives from John von Neuman's work. The content-addressable memory was studied, in connection with human memory, by William James.** Undoubtedly, both arrangement by position and arrangement by content have their merits, but each is suitable for different tasks.

"Location number" in a way is one attribute of stored data. But it is an externally assigned one, an attribute that is not necessarily shared by the same data existing in another memory. This makes the Neuman approach quite inflexible and less and less suitable as the complexity increases. Feigenbaum concurs with this and states that "Experience has demonstrated that the [von Neuman] arrangement limits extremely the behavior

---

*image will be used here to mean an item (data, records, instructions, etc.) that is stored in memory.

**James, W., The Principles of Psychology. New York: Holt, 1890.

of complex programs. The larger the computer memories get, and the more involved memory address housekeeping becomes, the more burdensome will become such selection system."*

Problems that require "intelligence" for their solution are often fairly complex. More important than that the solution path is ill-defined, ill-structured. The final solution often is path-dependent, i.e., is a function of what relations have been perceived and what have not. This puts a high premium on innovation (making new and plausible combinations of known relationships and facts) and this clearly requires the facility to pool together related facts and relationships. Such pooling together could be done with much greater ease if retrieval were possible using any one of the attributes of data (one attribute of which may be its position in memory), i.e., content-addressability.**

(2) Parallel search or simultaneous access. Given the desirability of content-addressability, we must decide which of the several methods possible to use.

One method is the brute force approach: serially examine all the memory cells. Obviously, as the memory gets large, this would prove quite impossible.

---

*Feigenbaum, E.A., An Information Processing Theory of Verbal Learning, RAND Memorandum P-1817.

**New computer scheme that achieves complex retrieval on the basis of content will be discussed in a forthcoming paper by the author, "Intermodular Communication Without Addressing in Planar Arrays of Modules Connected With One-way Channels", MIT, Sloan School Working Paper.

Another approach is to bring a "buffer" program between the user and the system so that although the system operates using location numbers for retrieval, the user needs to specify contents or attributes only. Such buffering has been realized in list processing languages like LISP, ILP, COMIT and SLIP. These languages "manipulate" symbols and attain content addressability at the expense of using excessive memory space.

For specific problems or tasks, LP's have proven useful. But there are several drawbacks. First, the data in the memory has to be quite structured, i.e., links or associations have to be explicit. Creating these links (although internally done) can be quite a chore as the data base gets large. Secondly, this as well as the first procedure are basically serial, i.e., search is done serially. This makes search time dependent on memory size. Therefore, for an information system with a data base that grows from period to period, continuous use of LP's may prove very cumbersome and impractical.

A third approach is achieving content-addressability by parallel search. Contents of either the entire memory or some portion of it are searched simultaneously. This facilitates very easy retrieval and storage. Such memory organizations have been realized and are reported as "associative memories" in the literature.

We feel this third approach would be particularly suitable when it is expected that the data base will be quite large. So because handling multitude of requests on the data base that would arise when the information system is imparted some intelligence would be very difficult unless scanning the entire memory (or suitable domains of it) is a fairly easy task.

It should be noted here that such an organization of memory would be very suitable for non-structured, path-dependent search; very tenuous hypotheses could be formed and easily linked with relevant information in memory through what corresponds to exhaustive search. This would allow one to grope around until concrete relations are found. It should also be noted that the memory merely makes "groping" feasible. Actual ability for this would derive from algorithms and heuristics in the system.

(3) Distribution of logic. There is one more question we should address ourselves to: How should the control be in the content-addressable, parallel search information system? Should the control be confined to one center (as is the case in many conventional computerized systems) or should there be several control centers? If the latter, what should the precedence scheme (importance hierarchy) be? This is nothing more than the decentralization-centralization issue that arises in the theory of firms and many of the arguments provided therein are applicable herein.

Inasmuch as the tasks which an "intelligent" information system will encounter will be fairly complex, we feel some degree of decentralization would increase the efficiency. Memory schemes exhibiting some decentralization have already been reported.* Such systems, called distributed-logic associative memories, provide each storage element with limited logic capabilities.

In the future, as such arrangements become more sophisticated, extensive decentralization should be possible, with each memory element (module) being by itself a finite state computer.**

---

* An example is the ILLIAC IV machine.

** Since the writing of this paper the author has developed and is about to patent a computer design scheme that will allow complete decentralization.

This type of memory would render the system very flexible. The memory could be organized and reorganized according to patterns most suitable for particular tasks or problems. Obviously, memory could easily be partitioned into relevant domains with respect to a given task. Then only the relevant domain need be searched or processed.

Facilities called for in assertions one through three constitute the framework for what shall be called here an associative information system. On this framework we need to build an operational structure and then supply a context for the system as well as tasks. The context will be indicated in assertion five below. An operational structure and tasks will be discussed later on.

(4) Human information system. Let us for a brief while look at a particular class of information systems with which we are in intimate contact twenty-four hours a day, namely, our own information system.*

We receive inputs, process and store them, and produce output. And we have somehow developed, and this is a very interesting point which will be explored later, a consciousness or a dichotomous personality such that as various processes go on in our system, we have the ability to observe some of these processes and how they relate to us.

Assuming the role of the inner observer, i.e., via introspection, what can we perceive about the operations of our information system? We see that we can retrieve items by specifying any one of a number of their attributes. We initiate very complex searches and the "attributes" of the results are used to retrieve more items. These suggest that our memory is content-addressable. How about parallel search? If we examine our thought

---

* The ideas discussed in this section have been greatly expanded in Sahin, "An Investigation Into Non-directed Thought Process and Their Pertinence to Information Systems Possessing Intelligence", MIT, Sloan School Working Paper 375-69.

train we find that it consists of a sequence of images. A cue (i.e., a part of an image) is used to generate the next image. This can be done if the cue is sent to memory and items containing that cue are retrieved (example: we think of our "car in the street." The "street" brings tp our mind the image or the thought that "street cleaners congress is in April." Clearly "street" could have been associated with a multitude of items but only one is chosen. An evidence to this and also of parallel search is that presented with a question or inquiry, we can quickly tell if we have answers or items in memory which meet the question or inquiry--although without specifying the item ("We have a feeling we have it in our memory."). Considering that human memory is fairly large (thus ruling out exhaustive serial search) it is plausible to infer that such certainties can take place if there is some kind of parallel-like search at least in some domain.

The location of the "processor" or "consciousness" in the brain have occupied the minds of many people. Many experiments have been aimed at finding it. So far no success. In fact, the mind appears to function when chunks of the brain from divers segments are removed. These would indicate the lack of a single processor and would point to the existence of "distributed logic." Experiments on neurons and simple combinations of neurons show that they very likely have logic capabilities.*

Let us explicitly state the implications. In trying to design an "intelligent" management information system, we posited, on the grounds of efficiency, that such an information system should have the facility of (1) content-addressability (2) parallel-like search, and (3) distributed logic. An inspection of the human information system, which certainly

---

* Bullock, T.H., "Integrative Properties of Neuronal Tissue," in Russell, R.W. (Ed.), Frontiers in Physiological Research, New York and London: The Academic Press, 1966, pp. 5-21.

exhibits plenty of intelligent behavior, reveals that it very likely has content-addressability, etc. This brings to mind cybernetic claims that many similarities are to be found between various systems--organic or not. Such similarities among systems addressed to similar problems should be all the more striking. This situation, i.e., similarity of problems, typifies the "intelligent" information system desired for firms and realized in human beings. We are suggested of a rather important thought: if indeed in the design of management information systems, we are headed in the direction of organic information systems and we are satisfied with what the human information system achieves, then we can speed the process of progress by directly borrowing from human information systems. We will state this more explicitly in terms of an assertion and say: Many concepts and schemes utilized in the human mind can and should be transferred to the design of intelligent management information systems. Some ideas that have been generated on some aspect of human information systems will be presented elsewhere.

(5) Suitability to planning and control. We feel an intelligent information system equipped with the facilities indicated in assertions one through three would be most useful in the planning and control activities of the firm. In these activities, search is often path-dependent. Much innovative thinking is needed, i.e., new and novel combinations of known relations and/or facts. Also, from a limited amount of information, it is desired that many conditional hypotheses be generated and tested with respect to our experiences in the past. Especially for effective control it is desirable to scan a large data base while at the same time generating and testing hypotheses about underlying relationships. These capabilities will be within the reach of the type of information system we have been discussing here and hence planning and control as its appropriate context.

Some Qualifications and Remarks: Some may feel that conventional facilities would be adequate to achieve intelligence. As is eloquently stated by Ashby* merely achieving intelligence is not satisfactory; we need to have it in enough quantities. Having the facility to answer some set of questions about baseball games (cf. BASEBALL program) is quite interesting. But this would hardly do for a management information system. The question answering ability must have a much wider coverage. It is for this reason that we posited the above assertions in that we felt attainment of "enough" intelligence would be facilitated by having content-addressability, parallel search, and distributed logic. Enough is the key word.** Also, it should be hastily added that there is nothing wrong with hybrid systems where part of the system is organized conventionally and the other part according to some new schemes. Monitoring of highly routine tasks would hardly require an "intelligent" system. But modification of the routine might.

State of the Art

We have surveyed the published literature to find out what has been done by way of achieving content-addressability.*** The approaches used fall into two categories: hardware approach where the memory itself is organized according to new schemes (associative memories) and the software

---

* Ashby, W.R., "Keynote Address--The Next Ten Years" Computer and Information Sciences, Julius T. Tou and Richard H. Wilcox, editors, Spartan Books, Washington, D.C., 1964.

** In making comparisons, if two systems are qualitatively found similar on some attribute, often the erroneous conclusion is reached that the systems are quite alike. Consider a monkey and a man. Qualitatively, similarities are overwhelming. But yet it would hardly do to say that functionally monkey and man are quite the same. We feel the difference between intelligence and "enough" intelligence is of this kind.

*** More details on the state of the art will be found in Sahin, "Associative Memories and Associative Processing: A Review of the Literature," MIT, Sloan School Working Paper

approach where content-addressability with respect to the user is achieved by implementing list processing languages on conventional equipment (associative processing). Only the hardware approach achieves parallel-like search and distribution of logic.

Hardware realization of the facilities (which we claim are needed in a full-scale intelligent managerial information system) seem possible on a limited scale. Associated memories with capacities of several thousand words have been constructed. Lee and Paull* report a distributed logic associative memory with "intercommunication" between cells. The cells, in this construction, have two states: active and inactive. If at a particular instant, a cell is "active" then it takes part in the processing (e.g., retrieval). This dynamically divides memory into relevant domains. A novel feature of Lee and Paull scheme is that cells can intercommunicate with adjacent cells in the sense that they can alter both state ("active" or "inactive") and the content of an adjacent cell and conversely itself be altered. Notice this in a way implies decentralization of control.

Let us draw attention to the fact that in an associative memory, cells are all identical in the sense that no cell is unique. This makes the system insensitive to local malfunction--non-functioning cell is simply left out. The data contents of the cell would, of course, be lost unless a duplicate record existed elsewhere in memory. Also, putting in and taking out memory chunks is facilitated. One can simply plug in a block of memory (with data stored in it, if you will) or take out a block. Such capability would be of tremendous importance to an information system. It paves the way to utilizing experience pooled in

---

* Lee, C. and M. Paull (Bell Labs), "A Content Addressable Distributed Logic Memory with Applications to Information Retrieval." Proc IEEE, Vol. 52, p. 312(L), March, 1964.

subsystems, e.g., in a firm where data processing is not centralized, a branch office can pool its experience in operation into a summary write-up coded into a memory block and then send this block to central processor. When the block is plugged in the central processor, the experience of the branch office immediately becomes a part of the knowledge store of central processor (or central office) and constitutes a part of the search space.*

Again because of a nonuniqueness property of the cells, they can be constructed in big batches. Recently developed integrated circuit construction is a very suitable mass production technique. In this process, using microphotographs of circuit diagrams, circuits are grown on silicon wafers by successive etching and oxidizing.** On a wafer no bigger than a thumbnail, hundreds of circuits can be placed. These circuits consume little energy and are long-lived as well as reliable.

Another point to note is that in associative memories, sorting and retrieval become trivial tasks whereas in conventional equipment this by no means is so. In fact, from this springs major advantages of associative memories.

Given that facilities called for in an associative information system are feasible and supposing they will be available on a wide scale in the near future***, i.e., that we have a general purpose, associative memory with parallel search and distributed logic, how do we structure a managerial

---

* Clearly this can be done by reading in the data, too. But the point here is that the "experience base" can be altered with much ease.

** Hittinger, W. and M. Sparks, "Microelectronics," Scientific American, Vol. 213, No. 5, Nov. 1965, pp. 56-75.

*** Ashby predicts we will have "distributed memory" and logic as well as parallel search "within the next ten years." Ashby, W.R., "Keynote Address - The Next Ten Years" Computer and Information Sciences, Julius T. Tou and Richard H. Wilcox, editors, Spartan Books, Washington D.C., 1964.

information system around it? In the next section we shall present some thoughts on this.

Structure of an Associative Managerial Information System

For the purposes of the discussion, we shall divide the information system into four stages: input, output, processing, and system modification. Each stage will be looked at separately. Eventually a unified model will be presented.

Input

The most important input decisions are: what activity to look at, in what depth and for what duration of time, i.e., what, how much, for how long.

Obviously, what we want to look at any moment of time is a function of what we can look at, i.e., our choice is dependent on the alternatives we face at any moment of time; but also on importance, according to some criterion, of each alternative. How do we determine importance? A criterion of importance frequently employed is the degree of expectability of an event. If we expect an event to occur and we witness it occurring, we often do not have much motivation to obtain more data on that event. Another criterion is the relevance of an event or activity to what we are doing, that is to say, to the state of the system. In other words, if what the firm is doing is tied to an event or activity, even though the outcome is highly expected, that activity or event may still be closely watched. Put another way, "importance" is determined by subjective priors about outcomes of events or activities and the expected payoff (or punishment) associated with each outcome.

After we have decided on what to focus on, then we have to start collecting information on this activity. These suggest that we need two setups or two networks: one for collecting information on a specific object or a specific activity and second one to monitor the first network. The second network would essentially decide what to focus on, for how long, and in what depth.

These two networks in an informal way exist in organizations, e.g., the accounting department collects information on particular events on a routine basis. The president, vice president, etc., observe certain changes (or new events) in the environment and the accountant is asked to extend his focus to the changes (or to new events).

The proposal here is that the above be done on a formal, explicit manner,using an associative information system. How can this formalization be achieved? To do this, we really need two distinct networks: one network having the capability of collecting detailed data on an activity or event, i.e., a focusing structure and a second network that does not collect detailed information but samples the entire environment and notes the major events and important changes in events.

The output of the second system which will be termed here the sampler is quickly relayed to the central processor and associated with questions, commands, and statements (facts) that exist in the memory. Note that the events in the environment are represented in the input stage by their attributes. So it is these attributes that are matched with items in the memory.

Let us stress that we are not only associating the external world with the facts that are contained in the memory, but also with instructions, so to speak, with "wants and desires." So incoming images may trigger instructions. Not only instructions but also questions.

A physical analogy may illustrate the point. Consider a radar base with a radar and a telescope. The radar corresponds to what we have called here a "sampler." It samples the entire environment namely the sky each time interval. The image generated this way is processed by the operator who processes the image associatively using his personalistic information system. In his head are stored statements (facts) which help him interpret the image (e.g., an object moving x squares on the screen in one time interval has a speed of y mph), questions (e.g., one "attribute" of the moving object might be that it is "unfamiliar." This causes the recall of the question "are there unfamiliar objects moving in restricted areas?") and finally instructions (e.g., the attribute "unfamiliar" may cause recall of the instruction "report unfamiliar objects in restricted areas to headquarter and also direct the telescope to them for more information"). The point here is that the radar serves to sample the environment. If events of interest are noted, the telescope is directed to the portion of the sky which contains the event. In this sense, the telescope corresponds to our analyzer.

Note that this sampler-analyzer system is an information system itself but in a well-defined environment.

Our interest is to achieve this capability in the information system of the firm.

As noted earlier, this dual-input associative information system in a way exists in the firm. The key executives receive summary reports. If they note changes (significant variances), they redirect the information system. Either the information system or a part of it focuses on the desired events and collects more data.

In a formal, i.e., explicit and well-formulated dual-input associative managerial information system, summary reports (from the samplers) would first enter the data base, perhaps on almost a real-time basis. Here associations would take place, relevant statements, commands,* and questions being retrieved. Unconditional commands would be executed or instructions for execution sent, and conditional commands would be presented to responsible decision makers. Among the set of commands might be some that pertain to change in focuses. These commands then modify the focusing activity.

Clearly, we have to devise some procedures and rules (or heuristics) whereby general "pictures" (i.e., summary reports of the environment) can be generated. Also, we have to do the same for analyzing and interpreting the general picture so as to distill the important attributes. These are not easy tasks but nor are they impossible.

Let us consider the problem of summarizing the behavior of the environment. First, we should define the environment which is of relevance (with regard to the firm). Then we should identify the eventspaces in this environment, i.e., divide it into sections and subsections, and sample each of the eventspaces separately.

Particular activities require particular measurement means (e.g., questionnaire, verbal) and methods (e.g., direct or indirect). Production activity is amenable to physical count, whereas customer satisfaction is not. Profit can be expressed in dollars and cents but is unobservable - we take raw data and calculate profit as a residue.

---

* We can divide commands into two classes: conditional and unconditional (voluntary and involuntary or reflexive and nonreflexive). The conditional ones are executed only after the responsible manager approves or perhaps modifies them.

These suggest that for each eventspace we should provide several means and methods of measurement, i.e., several types of "sensors" and also provide for mapping of signals into events, that is to say, consolidate input from several sources into an image of the underlying event. Let us once again examine the human information system. We find five types of senses that detect perturbations in various kinds of energy. The eyes deal with perturbations in light. The ears sense changes in the kinetic energy or air molecules and so on. So for the relevant environment, appropriate sensors have been provided. Without artificial aids, we don't perceive changes in radioactivity. This remains outside the relevant environment.

All of these sensors have elaborate transducers that convert the particular external energy perturbation to one type of perturbation - electrical. Within the nervous system very likely only coded electrical signals are used. These signals carry information about the stimuli and also the sensor - we know what we saw and what we heard. But more important than this, is that the codes (if there are distinct codes for each sensor) are easily translated (or mapped) into other codes. We know that the signal corresponding to the sight of the train refer to the same object. Sometimes such associations are so strong that the sight of an object evokes the auditory stimulus, e.g., upon seeing a train one "mentally" hears the rumble of it and conversely the rumble of a train brings to mind the visual image of a train. The mind with great ease performs this equivalence relationship.

In a firm there too are many sensors. And an event may manifest itself to several of these sensors at a time. In trying to design a "sampler" then, we face the problem of consolidating input from several sources into an image of the underlying events.

Output

During recent years, "selective dissemination of information" has become more popular.* In this type of output, the managers (or information receivers) have "profiles" stored in memory which indicate what types of reports the person wants to and should receive. Requests for reports not specified in the profiles are handled directly. However, the frequency and nature of requests are used by the information system to modify the profiles. If a manager puts a request for a special report rather often, after a while he starts getting that report automatically and regularly.

As it is now, this process is done in a serial and predetermined way in that the profiles are scanned in some definite sequence and according to a schedule.

We feel the profiles should be in the memory like any other piece of data and hence form a part of the domain with which items from the input channels or items resulting from activities in the processor are associated. When associations are found, the profiles and the associated items are processed together.

Such an approach has some noteworthy implications: (a) The distribution of reports is more tied to the activities of the information system. Inasmuch as the profiles form a part of the active search domain, the results of what goes on in the information system as well as input coming in are associated with the profiles; (b) By matching the image from the sampler with the profiles on a continuous basis, the input decisions (i.e., monitoring of the focusing network) would be much more efficient with respect to providing managers with relevant data; (c) If several profiles associate with a particular

---

* Hensley, C.B., "Selective Dissemination of Information (SDI)State of the Art in May 1963), AFIPS Conference Proceedings, 1963 Spring Joint Computer Conference, Vol. 23, pp. 257-263.

report, this would form the basis of inferring similarities and patterns among profiles; (d) To a significant degree, the decisions made and activities undertaken by managers are influenced by what information they get. Same report given to two different people in the organization will very likely evoke different types of behavior. Sometimes it is hard to know whose behavior will be appropriate for a problem at hand. The point is this: The pattern of information dissemination is an important contributing factor to the pattern of organizational behavior; for some problems, different organizational behavior and hence different information dissemination patterns may be more effective; deciding on a particular dissemination pattern would involve looking at profiles and also at relevant history of managers, generating several possible patterns and finally evaluating these patterns to choose one of them. These activities require scanning a large data base and forming many associations, as well as hypotheses. An associative information system, therefore, would be quite suitable for bringing about associatively selective information dissemination in the firm and hence precipitating partial modification of organizational behavior.

As in the input stage, in the output stage too we have to (1) decide what to put out, in what detail, when, and to whom and (2) carry out the output decisions. Paralleling this, we need two output networks, one to sample the internal output requests and a second to carry out chosen output requests. The first network would look at (a) special output requests, (b) profiles (according to priority schemes) (c) the activities of the processor (d) incoming data(e) some sets of instructions, and (f) at a set of hypotheses or conditional statements. The second output network is then monitored to carry out the information dissemination.

This dual-network output scheme would allow time-changing focusing for the output function, this being analogous for the focusing in the input function.

Processor

The contents of the memory can be classified into three types: (1) statements (raw data, records of past events and activities, policies, beliefs,* strategies, hypotheses, theorems, conditional statements, and facts), (2) questions (internal questions which constitute the tests applied to data, and external questions), and (3) instructions (internal instructions used for transforming data, and external instructions. Either set may contain conditional instructions).

The memory itself might be of three kinds: (1) active - long-term, (2) active - short-term, (3) potential (or peripheral). More will be said about short-term below.

We will also distinguish between two kinds of processing decisions: reflexive and nonreflexive (or compulsory and optional). The processing itself pertains to input activities, output activities, and to a set of functions to be performed (some of these are inductive inference, problem finding and evidence generation. These will be discussed in the next section).

Processing consists of two cycles. During cycle one, processing requests are "associatively" scanned, appropriate ones being chosen. During cycle two, the chosen processing activities are carried out.

A "follower" equipped with a short-term memory keeps a brief account of the processing activities. Such a follower is needed for creation of a model of the "self" essential for system modification.

---

* Here belief is used to mean a hypothesis that involves a great number of events but that has not (or cannot) be conclusively proven or disproven.

System Modification

How complex is complex? Ross Ashby provides an example of how quickly complexity can get out of hand.* Consider a block of lamps, twenty by twenty. A lamp is either "on" or "off." If each combination represents a pattern, then there are 2^{400} patterns. Supposing we define a binary property on each pattern, e.g., a pattern is either acceptable or unacceptable. This yields $2^{2^{400}}$ properties. If we observe that $10^{10} - 10^{8} \approx 10^{10}$ we realize how large this number is. For more concreteness, let us divide this number by the number of atoms in the visible universe, which is 10^{73}. We get $\frac{2^{10^{180}}}{2^{220}} \approx 2^{10^{180}}$. The number is not even touched.

The example dramatizes the fact that the complexity of our environment is insurmountable by brute force no matter how "large" our computation systems get. What we need are ways of reducing the complexity by many orders of magnitude, through learning more about the environment and discovering some order, some "formula." This would be even more so for an associative information system where the complexity increases exponentially. Clearly, we need to have in the system schemes for learning and schemes for learning to learn. Such schemes would require at least the following facilities:

(1) Generation of hypotheses about functional relationship (e.g., $y = f(x)$) and about logical relationships (e.g., $y \rightarrow x$ or $y \varepsilon x$). Induction schemes, to be explained a little later, could constitute a source of hypotheses.

(2a) Testing. The system should have the facility to test hypotheses both against internally stored data and against fresh observations obtained for the purpose. The output network would also be manipulated to test

---

* Ashby, Ross, "Keynote Address--The Next Ten Years", Computer and Information Sciences, Julius T. Tou and Richard H. Wilcox, editors, Spartan Books, Washington, D. C., 1964.

hypotheses pertaining to action-reaction relations, e.g., the impact of a particular information dissemination pattern. Results of testing would generate new hypotheses.

(2b) Experimenting. This is different from testing in the sense that one need not have a hypothesis but merely a question. Such activity would take place in response to an idea like "I wonder what would happen if I did..." This is analogous to mutation in genetics. Once in a while the system alters some part of its operations or strategies. If the results are favorable, all or some of the changes are retained.

(3) Evaluation. One needs to provide the system with facilities to perform evaluations which would form the bases for the decisions by the system to modify itself. It is our contention that to do this, the system first must be able to generate and maintain a model of itself, and also of its environment. Secondly, it must be provided with and find and also establish relationships between the model of itself and models of the environment. And finally, the system must have a structure of reference points, i.e., a value system so as to evaluate various relationships and events.

Inasmuch as the information system is to operate in and for a firm, it is desirable and even necessary that "its" value system be in many ways similar to that of the firm. How can this be done? Norbert Wiener suggests that values can be transferred to nonbiological systems the way they are transferred to biological systems, i.e., through exercises. For instance, a student presents his thoughts on a problem as part of an exercise. The relevant members of the society (his teachers) by imposing their own reference scheme on the answer, evaluate it and provide the student with

specific criticisms by which the student learns which of his thoughts were "wrong" or "unacceptable" and in the process acquires some of the teachers' references or values.

In a similar fashion, an associative information system would be assigned by the firm certain tasks or drills. The output would then be corrected and specific comments fed in. Note that part of the values acquired would pertain to "importance," i.e., priorities.

To make any use of the comments and bring about modification, a model of the self is quite essential. We have in mind here a dynamic or time-varying model. For the generation, and maintenance of dynamic self-models, the "follower" (described earlier) which is equipped with a short-term memory and which keeps a summary of the activities of the system would be indispensable.

We fell that in an information system which has a model of itself, models of the environment, and a value structure, and which can manipulate the models and evaluate the results, would evolve what could be called organizational "consciousness" or self-consciousness, for what is consciousness but an awareness of the self, awareness of alternatives, and awareness of how the alternatives and their outcomes relate to the self.

In the next section, we shall present some thoughts on several intelligent activities that would be performed in the processing stage of an associative information system.

Some Intelligent Tasks that can be Performed in an Associative Information System*

Insight (or Delayed Response)

It is a common experience that sometimes we find a speedy solution (obtain a "sudden insight") to a problem that had been bothering us for a long time. A possible mechanism to explain this phenomenon may be this: we work on a problem (or a question) but because we don't have in memory a needed fact or relationship, we cannot produce the answer. However, we have already done some processing of the matter. So we store the question with "tags" that it is unaswered and of such priority, and also store the amount of processing done on it. Some time later, when we come across the needed fact or relationship, the attributes of this may serve as the cue to retrieve the unanswered question and result in its being reprocessed. But since already some processing was done on it and since the needed data are now in memory, an answer is produced quickly. Note that a question may be recalled and reprocessed several times without being answered. However, each reprocessing brings the solution closer, so that when the last needed step is taken, the answer appears in a "flash".

Whether this actually happens in the human mind would be an interesting study in itself. However, we feel (1) such a scheme would be very suitable for implementation in an associative information system since it (the scheme) requires content-addressability (matching of attributes) and some parallel-search (pairing of incoming facts with stored items on a continuous basis

---

* A much more extensive discussion of these functions will be found in Sahin, "An Investigation into Non-Directed Thought Processes and Their Pertinence to Information Systems Possessing Intelligence," MIT, Sloan School Working Paper

would otherwise not be possible) and very possibly distributed logic, and (2) its implementation would be an important addition to the question-answering capability of the information system. A manager would address a question to the information system. If the answer cannot be provided, it would be stored in memory. Inasmuch as via parallel search, if not all at least some domains of memory are associated continuously with incoming data, the question is "active" all the time. Just parallel search would not be enough because to a given piece of data, far too many questions correspond. Distributed logic is needed to perform further tests at local level on data and questions stored.

Note that implied in this scheme is an associative output heuristic. When the long unanswered question is finally answered, it is presented for dissemination. The first output network compares the priority of this task with those of others. If priority is high enough, the output flow is altered so as to accommodate the execution of this task.

Such a scheme would eliminate to a great extent making up files of unfilled requests, and scanning them every so often and serially and then testing data base to see if relevant facts have come--an involved chore indeed.

This scheme would also greatly facilitate execution of standing requests of "reminders." Managers may request, for instance, that they be informed when such and such data comes or when such and such happens. Using current systems, executing a large number of reminders would be fairly difficult. In an associative information system, this operation could be part of the routine activities.

Evidence Generation

Often a manager gets some "hunches" or "suspicions" or "ideas." In trying to verify these, he talks to people around him, checks out some company files to see what the company experience was. But his search is necessarily limited. The search for evidence in the recorded experience of the company, to support (or contradict) the hunches, could be done more exhaustively and with greater ease in an associative information system. The hunch or hypothesis is divided into premises (if such and such) and consequences (then such and such). The attributes of the premises are used to retrieve records of events which contain the premises. Then the existence of attributes of the consequences in these records is tested. Should the consequences be found to hold in a majority (or some other criteria) of cases, the hunch is supported and some set of the supporting cases are presented to the manager.

Such a scheme would in general be used for hypotheses testing also.

Note that if we didn't have parallel search in the inforamtion system, scanning the "records of company experience" would be a highly time consuming task.

Inductive Inference

Inference can conceivably be obtained in a variety of ways. We propose the following scheme for implementation in an associative information system: As the attributes of an item are associated with items in memory which contain those attributes and these items are retrieved, some set of tests are applied to the items retrieved to find out if some attribute exists in them. Supposing that the search was on the basis of attribute x. Then as a result of the tests it was found that attribute y exists in most of the retrieved items, then it is inferred that y and x are related. Further testing would establish the connective between y and x, e.g., $y \rightarrow x$, or $y \subset x$. The inference

is then stored in memory. This way the inference becomes analogous to a "standing question" in that as fresh data pertinent to it comes, it is paired with them and tested anew (and also used in the processing).

The scheme can be implemented for generating inferences directly: A cue is selected and used to retrieve items containing it. Then the existence of other attributes is tested. This would be analogous to what a person does when he says, "I wonder what common characteristics (other attributes) divisions with million dollar sales (the cue) have?"

Again, we will note that inductive inference on a continuous basis would be quite difficult without content-addressability, some parallel search, and a degree of distributed logic.

Problem Finding

Professor Pounds of MIT argues that problem finding, which is a vital activity, does not occupy enough of our efforts.* He defines a problem as the difference between what is actual and what is desired. He finds as a result of his observations in a manufacturing company that historical models are often used by managers to find differences and hence problems. We feel the definition of the problem given by him while not all inclusive is quite useful. The procedure of problem finding can be readily implemented in an associative information system using historical models, "scientific models," extra-organizational models, etc. In fact, detecting differences between incoming reports and desired results on the basis of models stored in memory can be made a "regular" activity of the processing stage. The "value and priority systems" and the "self-model" discussed earlier would be used to select the noteworthy problems. The "associatively selective dissemination network" would then relay reports on the problems

---

* Pounds, William, "The Process of Problem Finding," MIT, Sloan School Working Paper 148-65.

to appropriate people. Notice that reply from a manager that a problem reported to him does not really constitute a problem would cause the associative information system to evaluate and modify its "value and priority" schemes.

Concluding Remarks

I have tried to sketch a framework for an associative information system. Clearly, many of the ideas presented need to be probed in depth. Also for operationality, much detailing needs to be done. A future study will aim to supply the details within the context of the planning and control activities of a firm. Ensuing this, implementing the system using the data base of a division or a department of an operating firm will be attempted.

BASEMENT

Date Due

APR 12 '78

JUL 2 1980

JUL 12 1985

APR 30 1990

NOV

Lib-26-67

MIT LIBRARIES
3 9080 003 874 655
370-69
MIT LIBRARIES
3 9080 003 874 465
371-69
MIT LIBRARIES
3 9080 003 874 663
MIT LIBRARIES
3 9080 003 905 244
MIT LIBRARIES
3 9080 003 875 280
MIT LIBRARIES
3 9080 003 875 298
MIT LIBRARIES
3 9080 003 906 317
MIT LIBRARIES
3 9080 003 906 358
MIT LIBRARIES
3 9080 003 875 306
397-69B
MIT LIBRARIES
3 9080 003 875 322